

ON-BOARD DEMUX/DEMOM

N 9 2 - 1 4 2 2 1

S. Sayegh, M. Kappes, J. Thomas, J. Snyder,
M. Eng, J. Poklemba, M. Steber, and G. House
COMSAT Laboratories
Clarksburg, Maryland 20871

Abstract

To make satellite channels cost competitive with optical cables, the use of small, inexpensive earth stations with reduced antenna size and high powered amplifier (HPA) power will be needed. This will necessitate the use of high e.i.r.p. and gain-to-noise temperature ratio (G/T) multibeam satellites. For a multibeam satellite, on-board switching is required in order to maintain the needed connectivity between beams. This switching function can be realized by either an receive frequency (RF) or a baseband unit. The baseband switching approach has the additional advantage of decoupling the up-link and down-link, thus enabling rate and format conversion as well as improving the link performance. A baseband switching satellite requires the demultiplexing and demodulation of the up-link carriers before they can be switched to their assigned down-link beams. This paper discusses principles of operation, design and implementation issues of such an on-board demultiplexer/demodulator (bulk demodulator) that was recently built at COMSAT Laboratories.

1. INTRODUCTION

A multiyear effort was undertaken at COMSAT Laboratories to investigate the on-board demultiplexer/demodulator concept to determine its feasibility, identify critical technologies, and assess the potential of developing these technologies to a level capable of supporting a practical, cost-effective on-board implementation. An important part of the effort was a review of the advances that can be expected to occur in the critical digital component areas in terms of power, mass, size, speed, and radiation resistivity of the digital, logic, and memory components from which the processor is to be fabricated.

A baseline system of the demultiplexer/demodulator was defined and its performance evaluated by analysis and computer simulations. A digital implementation was selected to provide the flexibility that permits the on-board processor to accommodate different types of multichannel frequency-division multiple access (FDMA) carriers simply by changing its computational rules and organization. This permits the rules and organization of each processor to be modified to accommodate variations in the number and bandwidths of carriers over the lifetime of the satellite or to accommodate different applications of the same type of satellite.

A block diagram of the overall system and test setup is shown in Figure 1. The system uses the frequency-domain filtering approach to demultiplexing and a shared high-speed coherent demodulator. The fast Fourier transform (FFT)-based demultiplexer is capable of processing a large number of carrier types and bit rates. The demultiplexer

output is fed into an interpolating filter whose task is to deliver 2 samples per symbol to a shared variable bit rate digital demodulator that operates on a number of different carriers in a round-robin fashion. The COMSAT digital processor performs demultiplexing/demodulation and associated filtering and control for a number of carriers occupying a bandwidth of 20 MHz. The architecture used in this system is very flexible, allowing in-orbit frequency plan reconfiguration under ground command.

Most of the hardware has been implemented in low-power complimentary metal-oxide semiconductor (CMOS) circuitry. Several other important developments contributed to very substantial reductions in the power, mass, and size of the processor. An application-specific integrated circuit (ASIC) gate array chip that performs the interstage reordering in the FFT pipeline was designed and developed. This contributed to better than an order of magnitude reduction in power and mass as compared with a discrete large-scale integration (LSI) implementation. A method for sharing a single pipeline inverse FFT processor among the different carriers was conceived. By interleaving the frequency samples of those carriers at the input to the inverse FFT (IFFT) processor and selectively bypassing butterfly operations, carriers of different bandwidths can be handled simultaneously in the shared pipeline. This obviates the need for a separate IFFT processor for each carrier. A novel PROM-based approach was implemented for the acquisition section of the shared digital demodulator, significantly reducing the required hardware.

The demultiplexer/demodulator presented above has been constructed and tested at COMSAT Laboratories and is now operational. System performance evaluation in terms of bit error rate measurements are presented in this paper.

2. FREQUENCY DOMAIN FILTERING

An FFT/IFFT frequency-domain filtering architecture was selected for the demultiplexer. FFT/IFFT frequency-domain filtering method basically consists of convolving the composite frequency multiplexed signal with a bank of filters using an overlap-and-save technique. It computes the desired linear convolutions in terms of circular convolutions. The circular convolutions are computed by transforming the time-domain quantities to be convolved to the frequency domain, multiplying the resulting frequency coefficients across the overall spectrum by any desired filter functions and transforming back to the time domain. Specifically, to obtain carrier k , the frequency multiplexed signal is transformed to the frequency domain by an FFT, multiplied by the frequency response of filter k (typically a square-root Nyquist that serves the double purpose of demultiplexing and matched filtering), and the product is then

*This paper is based on work performed at COMSAT Laboratories under the sponsorship of the Communications Satellite Corporation.

transformed back via an IFFT to recover the time-domain waveform. Therefore, to obtain the individual baseband signals, the number of inverse transforms to be performed equals the number of carriers N . To minimize the amount of computation involved, the IFFT performed on a given carrier should only cover the frequency band occupied by that carrier. Thus, different carrier bandwidths will result in different IFFT sizes.

Figure 2 summarizes the frequency-domain demultiplexing approach. An important consideration here is the size of the Fourier transform that has to be performed. If the filter impulse response has L coefficients and 50-percent overlap is used between blocks, then the size of the Fourier transform to be performed is $2L$. Note that as the overlap between blocks increases, so does the number of computations per output sample. On the other hand, if the overlap decreases, then the size of the Fourier transform, and hence the memory size, increases. A 50-percent overlap achieves an almost optimum tradeoff between computational and memory requirements and is very simple to implement.

A pipeline FFT processor is an efficient way of performing the needed high-speed, real-time Fourier transforms by distributing the processing among several computational elements. It has a compact and modular structure and is well suited for very large-scale integration (VLSI) implementation. The pipeline processor consists of two building blocks: butterfly computational elements and delay-switch-delays (DSDs). The computational elements perform the necessary butterfly computations of the Cooley Tukey FFT algorithm. The DSDs consist of shift registers first-in first-out (FIFOs) and switches. Their function is to present the samples to the butterfly computational elements at the right place at the right time.

A radix-2 or a radix-4 implementation may be used for the FFT/IFFT pipelines. Although the radix-4 butterfly computations are more involved than those of the radix-2 (three complex multiplications *vs* one for the radix 2), the number of butterfly stages is half that required for a radix 2 and they operate at half the speed. Therefore the number of complex multiplications per second is 25 percent smaller for a radix-4 implementation. The choice of radix for the FFT is thus a tradeoff between speed and additional hardware. If the speed requirement can be satisfied by either implementation, then radix 2 may be the preferred choice. At the time the proof-of-concept (POC) model was designed, the radix-4 pipeline, shown in Figure 3, was chosen because of speed considerations.

As more highly-integrated devices become available, however, this choice must be reconsidered. A radix-4 butterfly chip operating at one speed and a radix-2 butterfly chip operating at twice the speed can each handle the same data rate. Thus, the answer to which is best turns to such factors as package size and power consumption, with consideration of the fact that the radix-2 pipeline requires twice as many butterflies.

As signal processing chips advance beyond the basic butterfly operation, the additional functions they include and the means of controlling them must be considered to

determine whether competing devices are more or less desirable. For example, one manufacturer may offer on-chip coefficient memory while another may not. A third may have coefficient memory that requires more off-chip control to utilize for our application. Therefore, the best architecture depends upon the total board area and power consumption required to perform a complete function (such as an FFT) at a particular speed.

The FFT pipeline in the POC model is capable of accepting four complex input samples and providing four complex output samples during each 11.52-MHz clock period. Thus, a 256-point FFT is computed every 64 clock periods or 5.6 μ sec. By extending the length of the pipeline, a 1024-point FFT could be computed every 22 μ sec.

A reduction in these times by a factor of approximately two would be a desirable objective for the near future in order to double the maximum IF bandwidth to about 40 MHz (which corresponds to a pipeline data rate of 80×10^6 complex samples/sec.). The long-term goal is an additional factor-of-two improvement to permit direct processing of IF bandwidths as great as 80 MHz.

The DSD is one of two key elements used to implement the pipeline FFT and IFFT processors (the other being the butterfly arithmetic processor). Due to the complexity (large amount of hardware) of the circuit, it is more practical to implement it with ASIC technology. COMSAT Laboratories has developed this DSD ASIC chip as part of the demultiplexer/demodulator program. A detailed description of the COMSAT developed DSD is now presented.

To implement one complete DSD function, eight ASIC chips and a small amount of discrete logic integrated circuits (ICs) are needed. In the FFT processor, one complete DSD is used between two butterfly elements and its function is to reorder the samples in its input data streams appropriately for the butterfly that follows it. The reordering process is achieved by using two sets of delay elements and one set of switch elements. The first set of delay elements are used to shift the input streams appropriately in time, then the switch elements interchange the samples in a predetermined fashion, and finally the second set of delay elements are used to shift the samples back in time appropriately. The DSD ASIC is hardware-programmable for the particular stage of the FFT or IFFT processor in which it is used. Specifically, there are four possible configurations for the DSD, three of which are for radix-4 transforms and one for the radix-2 case used in the IFFT processor. In the radix-2 case, the DSD treats the four input streams as two groups of two input streams. The DSD configured for the stage closest to the output of the processors has the smallest amount of delays and the highest switching rate.

The functional block diagram of the DSD ASIC is shown in Figure 4. The number on the right side of the delay element blocks indicate the delay values associated with the particular input data stream. For example, the four inputs X11-X14 always have delay values of 0, the four inputs X21-X24 can have delay value of 1, 4, 8 or 16 clock cycles. For any one configuration selected, only one set of

delay values are used for the input and output data streams. For example, if the DSD is used in the second to the last stage of the FFT processor, the delay values 0, 4, 8 and 12 are used. Specifically, the signals X11-X14 and Y41-Y44 assume the delay value of 0, the signals X21-X24 and Y31-Y34 assume the delay value of 4, the signals X31-X34 and Y21-Y24 assume the delay value of 8, and so on.

The switch elements in the DSD perform the function of routing the incoming signals to the appropriate outputs of the switch elements. With reference to Figure 5, there are two configurations for the switch elements, one of them is the 2-state and the other is the 4-state case, and they are used for the radix-2 and radix-4 applications, respectively. For the 2-state case, in state '0', the inputs to the DSD go straight through it, and in state '1', data streams at ports INA and INB are interchanged and data streams at ports INC and IND are interchanged. For the 4-state case, the situation is slightly more complicated, and the actions taken by switch elements in each of the four states are shown in Figure 5. For any one of the two switch configurations selected, the switch elements always go through the same states. The only difference when the DSD is used in different stage of the FFT or IFFT processors is the rate at which the states are switched. Specifically, the DSD closest to the FFT processor output has the highest switching rate, and it switches state every clock cycle. The DSD in the preceding stage switches states once every four clock cycles, and so on.

With reference to Figure 6 and the functional details of the DSD ASIC mentioned above, the implementation description is presented next. The delay elements are implemented using shift registers and 4-to-1 multiplexers. For a particular data stream, the possible delays of the data is achieved by connecting the outputs of the shift registers from the appropriate output stages to the inputs of the associated 4-to-1 multiplexer, and by selecting one of the four inputs as the output.

The switch elements are also implemented using 4-to-1 multiplexers, and by selecting the appropriate inputs, the data interchange as indicated in Figure 5 can be accomplished. The controller of the DSD ASIC is responsible for providing all the timing and control signals for the shifting and multiplexing operations within the ASIC.

Whereas a single FFT is performed on the composite frequency-division multiplexing (FDM) signal, the IFFTs are performed on an individual carrier basis. The case of mixed carrier sizes (narrowband, and wideband) is readily handled by performing inverse transforms of larger sizes for the wideband carriers. In order to avoid the duplication of hardware, a common pipeline, capable of performing transforms of various sizes under software control without the need for physically adding or removing any modules, is desirable.

When a pipeline is dedicated to performing an FFT of a given size, the number of stages in the pipeline is fixed and the twiddle factors of the butterfly computations, as well as the switching times and delays of the DSDs are readily available. The pipeline can be modified to perform transforms of various sizes simultaneously. The needed modifi-

cations are performed dynamically (using a few control signals) to allow the pipeline to constantly alter its function in real-time to accommodate the various transformation sizes required. By properly ordering the input data to the pipeline and bypassing some arithmetic modules for the smaller size transforms, any mixture of IFFTs whose sizes are a power of the pipeline's radix can be performed without requiring any changes to the simple and regular action of the DSD, as illustrated in Figure 7.

3. INTERPOLATION

Because the FDMA signal consists of asynchronous carrier transmissions, the samples at the demultiplexer output must be interpolated before being presented to the demodulator. The interpolating filter module (IFM) which connects the demultiplexer output to the demodulator input performs two functions. First, it adjusts the number of samples per symbol for each carrier from an arbitrary value near two to **exactly** two. Second, it adjusts the sampling point for each carrier to coincide with the peaks and zero crossings of the signal. It performs both functions by means of adjusting the phase shift of a simple finite impulse response (FIR) digital filter. The control signals for adjusting the number of samples per symbol are generated locally and asynchronously and are added to the accumulated clock error fed back from the demod to produce a composite control signal proportional to the instantaneous phase adjustment for the current sample. This signal is fed to the FIR filter.

Figure 8 shows the shared control circuitry of the IFM. The upper part of the diagram shows the circuitry required to generate the coefficient programmable read-only memory (PROM) addresses as well as general control signals. Each carrier address counter keeps track of the location within the current phase plan used for correcting the number of samples per symbol. This address is fed to a phase plan lookup PROM for each type of carrier in use. These PROMs are shared by different carriers of the same bit rate and coding scheme. This signal is then added to the output of the clock error accumulator for each channel and applied to the coefficient lookup PROMs to obtain the coefficients for the FIR filter. As a practical matter, the coefficient PROMs shown are duplicated on the second board to avoid too many board-to-board connections.

Figure 9 shows the nonshared circuitry of the IFM, i.e., circuitry that must be repeated for the I and Q channel. The shift register, multipliers, and adders constitute the basic FIR filter circuit. The data buffer and related control circuitry provide samples on demand to the input of the FIR filter. This is necessary in cases where the samples in the shift register must be reused to generate two output samples. In other cases where the current contents of the shift register are not required for an output sample the outputs are simply marked as invalid. This peculiarity occurs due to the fact that the number of input samples does not match the number of output samples.

In its current configuration, the entire IFM occupies two 9Ux440 wirewrap cards and uses predominately high-speed

CMOS digital logic devices including the LSI Logic L64012 multiplier IC and the Logic L4C381 accumulator IC. The first board contains the shared control circuitry and the I channel filter while the second board contains the Q channel filter.

4. DEMODULATOR

The on-board demodulator operates on a multiple set of quadrature phase shift keying (QPSK) modulated, asynchronous carriers in a TDM format, where the incoming TDM data packets are typically a fraction of the transmitted burst. In this manner, the demodulator processes only a few symbols for a given carrier, stores the results, and preloads its registers with the appropriate sample values for the upcoming carrier. The sample rate entered into the demodulator for all carriers is two samples/symbol. Recall that the sample frequency for all carriers is the same as their symbol rates after they have been warped in the FDM/TDM conversion. Symbol timing feedback from the tracking loop to the preceding interpolating filter places the two samples into the demodulator at the data-detection and symbol-transition points. The receive Nyquist data shaping has already been done in the receive filter module. However, the sample values at the data-detection points are modulated by a beat note between the actual incoming center frequency and the front-end down-conversion local oscillator. A carrier-phase rotator, which is effectively a 2×2 matrix multiplication of the beat modulated I and Q channels, is employed to remove the beat as follows:

$$\begin{bmatrix} I_k \\ Q_k \end{bmatrix} = \begin{bmatrix} \cos(\hat{\theta}) & -\sin(\hat{\theta}) \\ \sin(\hat{\theta}) & \cos(\hat{\theta}) \end{bmatrix} \begin{bmatrix} I'_k \\ Q'_k \end{bmatrix}$$

where $\hat{\theta}$ is the carrier tracking loop output phase estimate.

With a "0101" acquisition preamble in both channels there is a potential 180° ambiguity in the recovered carrier phase, which is resolved by means of the unique word (UW). The UW pattern is the same in both channels, so binary decisions used to increase detection reliability.

There are two phase-locked tracking loops in the demodulator for carrier phase and symbol timing. The carrier-phase tracking is second order to account for frequency offsets, whereas the symbol timing is first order and only tracks slow-varying phases. Multiplier-accumulators (MACs) are used to implement the digital tracking loops. The accumulators are preloaded with initial-phase or frequency information, whichever is appropriate, from the acquisition estimator circuitry. In this manner, the phase-locked tracking loop synchronization in burst mode can be expedited and more reliable. In terms of the second order loop parameters, the phase and frequency multiplier gains for carrier tracking are selected respectively as

$$K_\theta = 2\zeta W_n T_s$$

$$K_{\Delta\theta} = (W_n T_s)^2$$

where ζ is the damping ratio, W_n is the natural frequency, and T_s is one symbol time interval. For the first order symbol timing loop the multiplier gain is

$$K_\tau = (W_n T_s)$$

Initial estimates for carrier phase and frequency as well as symbol timing are computed in the acquisition estimate processor as shown in Figure 10, and briefly described as follows. Incoming I and Q channel samples are multiplied by a bipolar alternating sequence to remove the preamble modulation and averaged to improve their signal-to-noise (S/N) ratio. This yields four outputs, namely, even and odd sums in both the I and Q channels. The sums are taken twice, over the first and second halves of the preamble. The carrier-phase error may be found from

$$\hat{\theta} = \tan^{-1} \left(\sqrt{\frac{Q_e^2 + Q_o^2}{I_e^2 + I_o^2}} \right) - 45^\circ \operatorname{sgn}(I_e Q_e + I_o Q_o)$$

Similarly, the symbol timing error can be related as

$$\hat{\tau} = \frac{2}{\pi R_s} \left| \tan^{-1} \left(\sqrt{\frac{I_o^2 + Q_o^2}{I_e^2 + Q_e^2}} \right) - 45^\circ [1 - \operatorname{sgn}(I_e I_o + Q_e Q_o)] \right|$$

In both cases, the primary estimate can be found from a lookup table of the inverse tangent of a ratio of squares, and the phase ambiguity can be determined from looking up the sign of a sum of products. Since these computations are only required at a rate of twice per preamble, common processing elements can be used where the differences between phase and timing are incorporated into the final value lookup tables. The final value of the carrier phase at the end of the preamble is

$$\hat{\theta}_0 = \hat{\theta}_2 + \left(\frac{\hat{\theta}_2 - \hat{\theta}_1}{P/2} \right) \cdot (P/4), \text{ modulo } 180^\circ$$

where $P/2$ is half the preamble length.

The carrier frequency offset estimate is determined as

$$\hat{w}_o = \frac{\hat{\theta}_2 - \hat{\theta}_1}{P/2}$$

Lastly, the timing estimates are averaged as

$$\hat{\tau}_o = (\hat{\tau}_1 + \hat{\tau}_2)/2$$

5. PERFORMANCE

The BER performance of the on-board demultiplexer/demodulator processor has been measured using the setup shown in Figure 1. Four modulators are used on the transmit side to generate FDMA/TDMA test signals. All four of the modulators are capable of variable bit rate operation and have synthesized carriers so that a wide variety of frequency plans can be generated. The fourth modulator can

be used as an interfering burst for TDMA measurements. Noise is added at the 140-MHz IF to the combined modulator signals before processing by the demux/demod. The BER of any one of the demodulated channels is measured by the performance monitor by comparing the incoming data with a stored version of the transmitted data. Synchronization is provided by the UW detect signal from the demodulator for the selected channel.

To evaluate the performance of the on-board processor carriers corresponding to 1.544 Mbit/s with rate 3/4 and 1/2 coding and 2.048 Mbit/s with 3/4 coding were utilized. As a baseline, the carriers were first processed individually providing single carrier performance. Next, all three carriers were generated and supplied to the processor, but the IFM was set up to process only one of the carriers. This selection effectively separates the demultiplexing and demodulation functions of the processor so that implementation degradations can be isolated to individual subsystems. Finally, all three signals were allowed to pass through the entire system with the BER monitor selecting one of the three signals. A summary of the performance for 1.544 Mbit/s carrier with rate 3/4 coding for the three setups is shown in Figure 11. As can be seen from this figure, there is a small amount of degradation when the three carriers are introduced relative to the single carrier performance, but very little additional degradation when all of the signals are being processed by the IFM and demodulator. This degradation is thought to be due to a slight nonlinear operation of the demultiplexer front-end and is being investigated. In addition, some flaring of the data occurs at the lower error rates for all of the curves resulting from low-level interference effects. Overall,

the BER performance data provides validation of both the overall demultiplexer/demodulator structure and the selections of bit resolutions made early in the program.

6. CONCLUSIONS AND SUMMARY

An architecture for implementing an on-board flexible demultiplexer/demodulator was presented. The architecture is based on a frequency domain filtering approach to demultiplexing an up-link FDMA signal consisting of a mixture of carriers of different bit rates was presented. Specially designed FFT pipeline processors were used for this purpose. An ASIC chip designed at COMSAT Laboratories as a critical part of the FFT/IFFT processor was described. A digital demodulator architecture that operates on the interpolated demultiplexer output was presented. A survey of current technology illustrated that for the near future high-speed low-power digital signal processing will be mainly based on Si technologies (CMOS and CMOS/silicon-on-sapphire [SOS]). Based on COMSAT's experience with POC developments of processors similar to the ones discussed in this paper, as well as projections of technology, it is estimated that an 80-MHz fully-digital, very-flexible flyable processor is an achievable goal for the late 1990s. Such a processor is projected to consume only 25 W and have a mass under 5 lb.

7. ACKNOWLEDGMENTS

The authors wish to acknowledge the support provided by S. J. Campanella and R. J. Fang in performing the work reported in this paper.

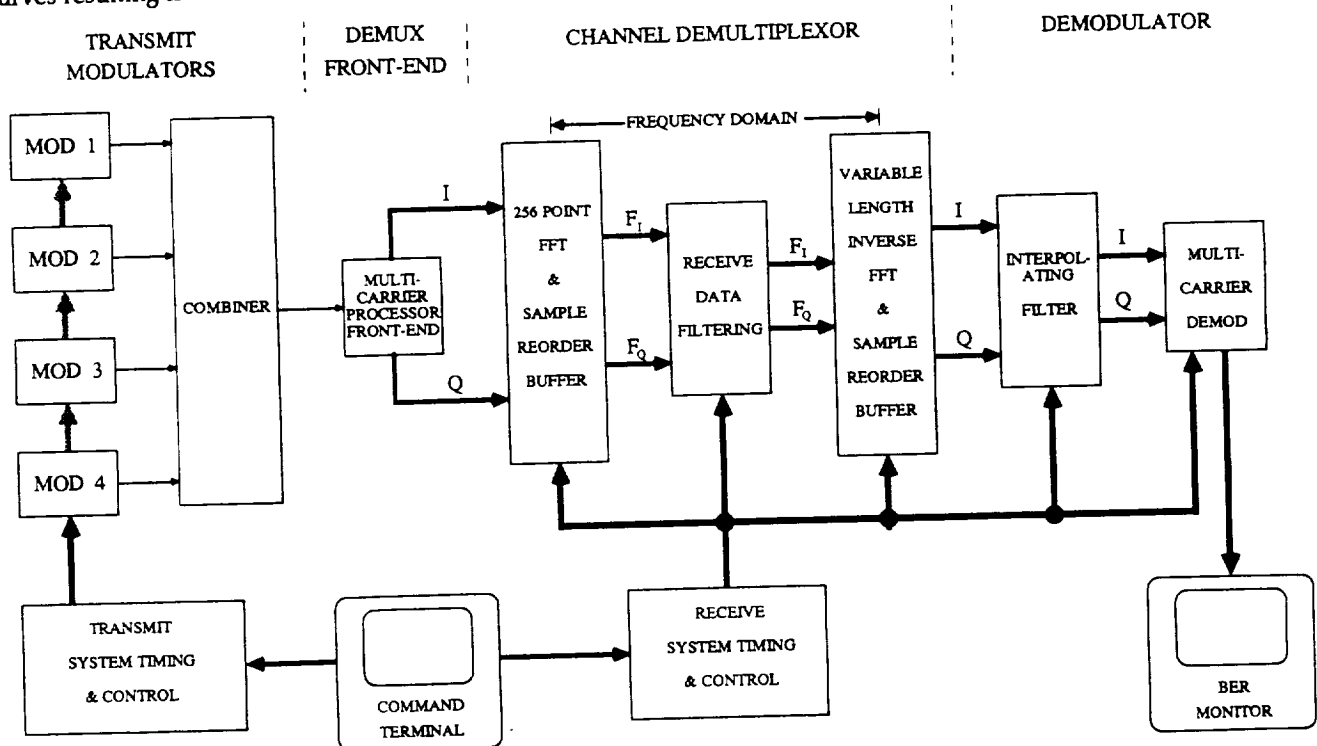


Figure 1. System Block Diagram

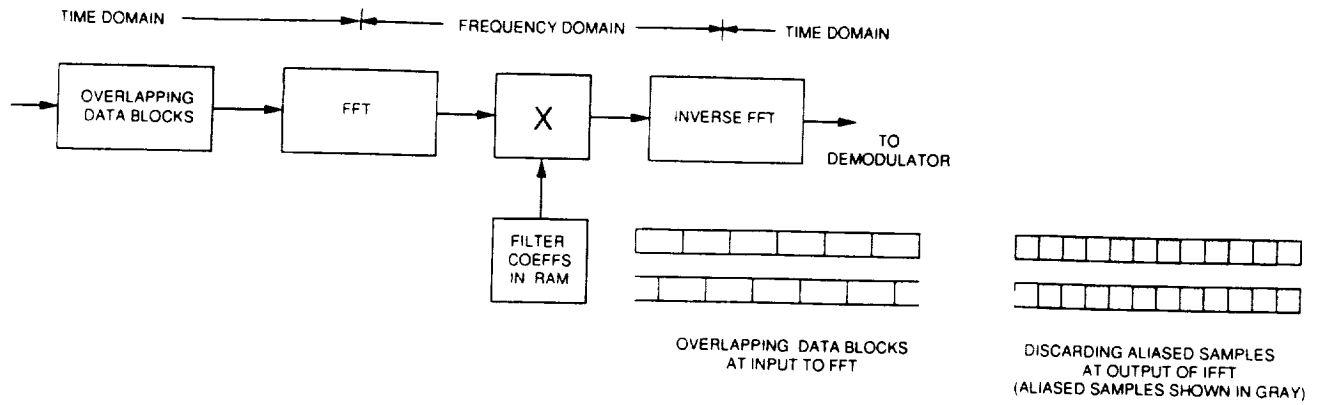


Figure 2. Frequency-Domain Filtering Approach

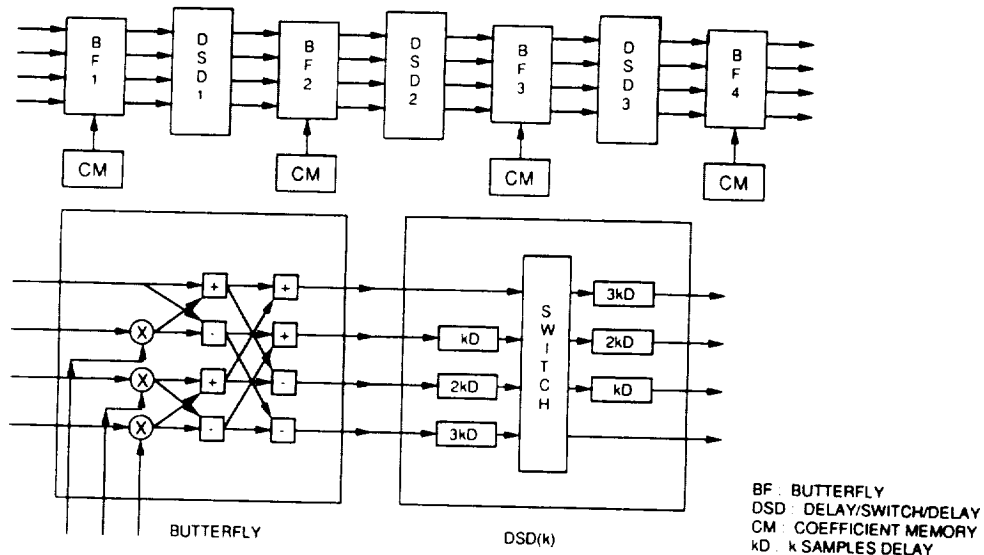


Figure 3. Four-Stage Radix 4 FFT Pipeline

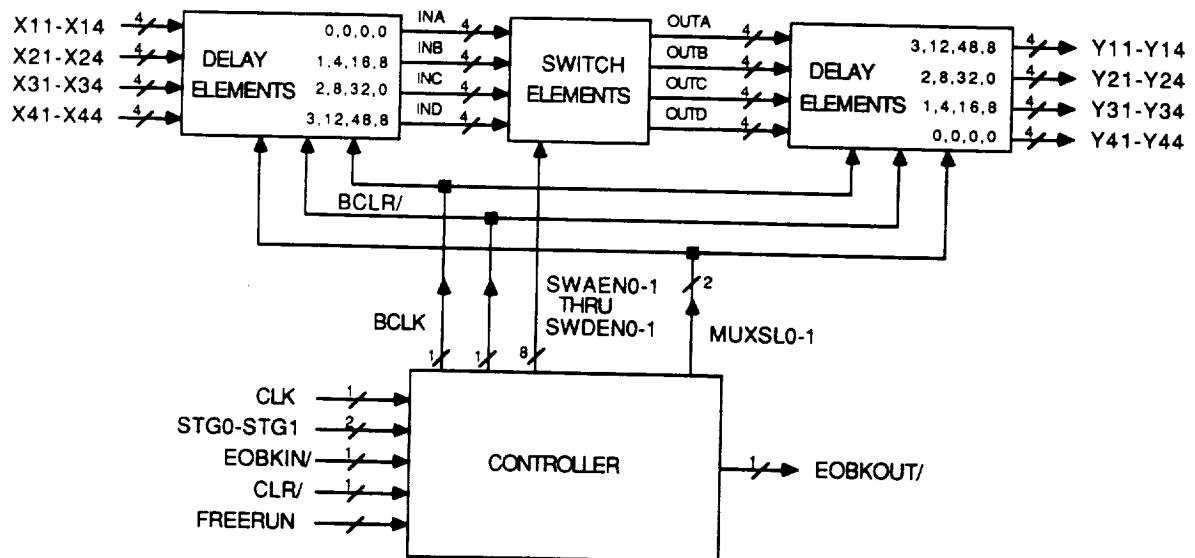


Figure 4. 4-Bit Wide Delay-Switch-Delay ASIC Functional Block Diagram




Figure 5. 4-Bit Wide Delay-Switch-Delay ASIC Switch State Diagram

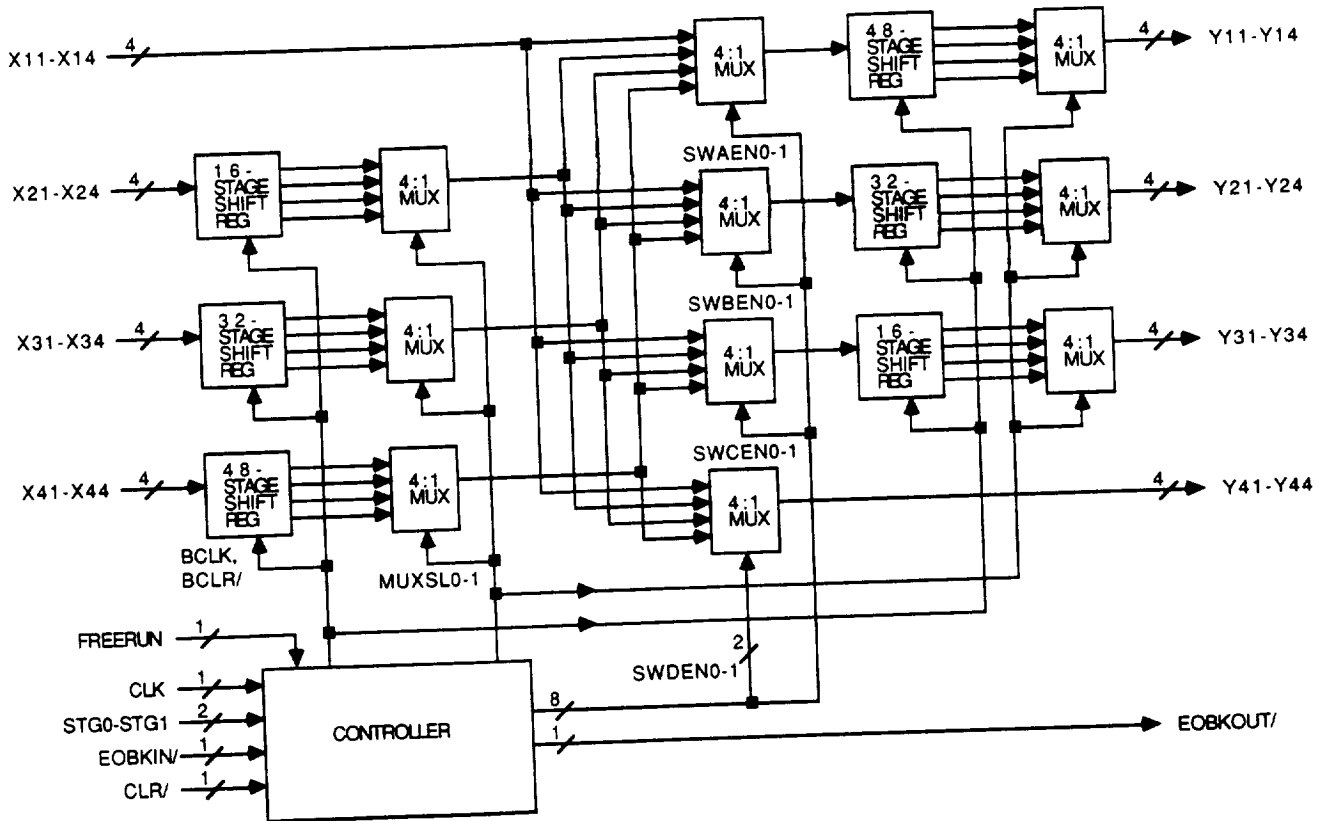


Figure 6. 4-Bit Wide Delay-Switch-Delay ASIC Implementation Block Diagram




Figure 7. Mixed-Size IFFT

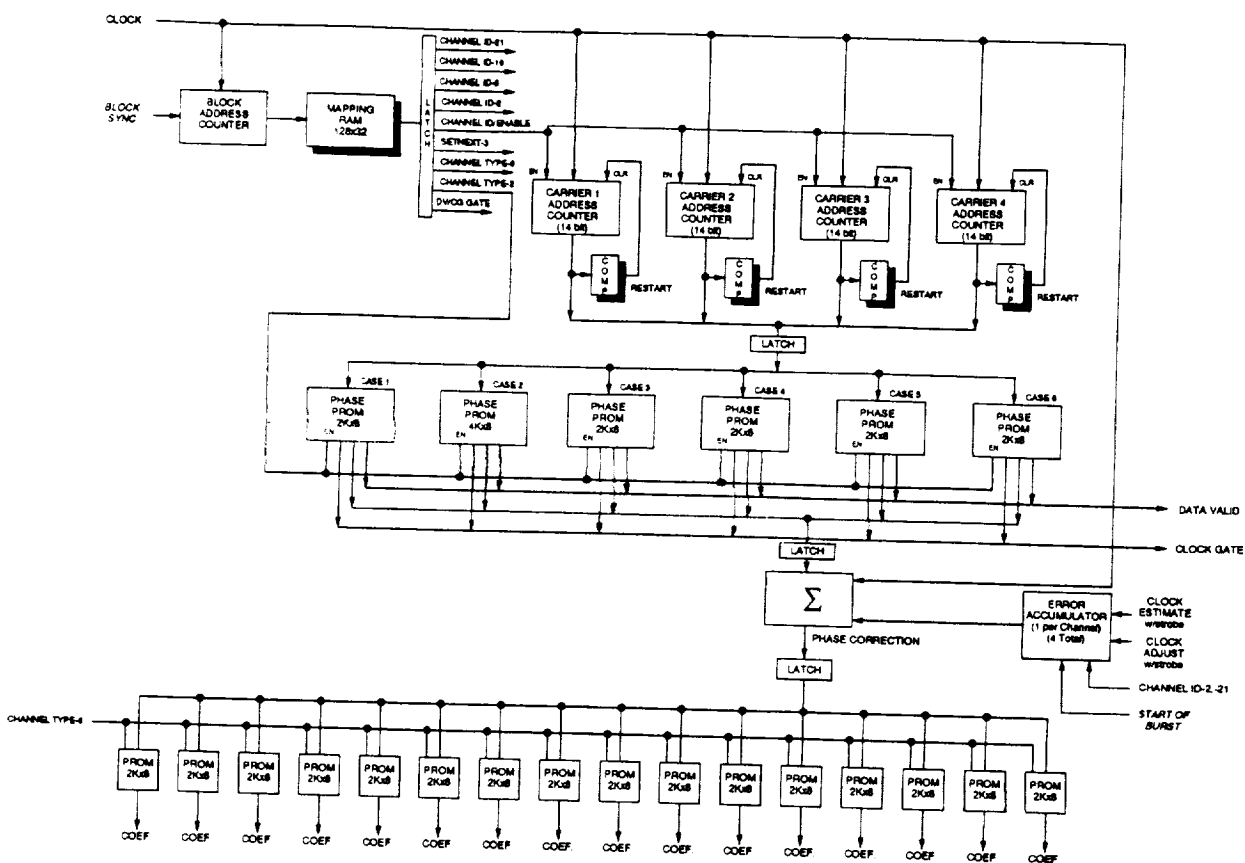


Figure 8. Interpolation Filter Control

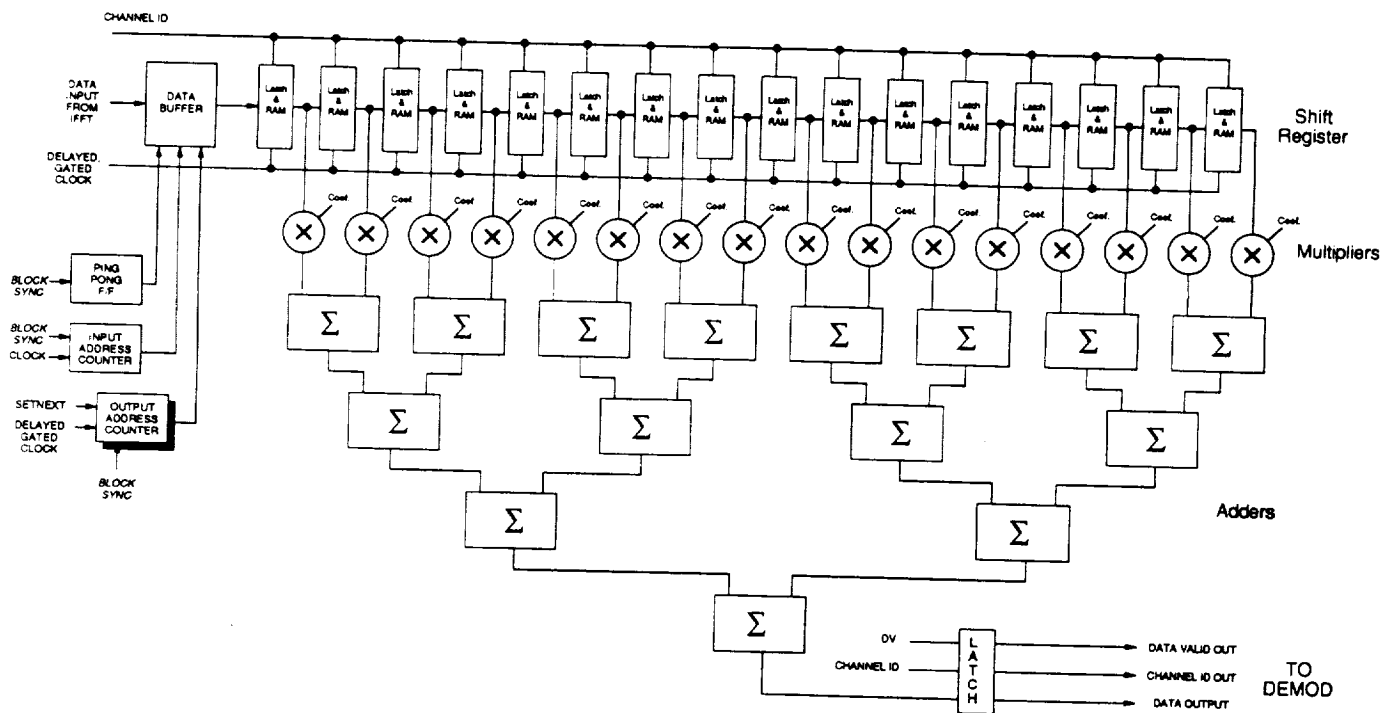


Figure 9. Interpolation Filter Computations




Figure 10. Demodulator Acquisition Control Module




Figure 11. BER Performance

